Hard-Clustering with Gaussian Mixture Models

Johannes Blömer, Sascha Brauer, and Kathrin Bujna

Department of Computer Science
Paderborn University
33102 Paderborn, Germany

March 22, 2016

arXiv:1603.06478v1 [cs.LG] 21 Mar 2016

Training the parameters of statistical models to describe a given data set is a central task in the field of data mining and machine learning. A very popular and powerful way of parameter estimation is the method of maximum likelihood estimation (MLE). Among the most widely used families of statistical models are mixture models, especially, mixtures of Gaussian distributions.

A popular hard-clustering variant of the MLE problem is the so-called complete-data maximum likelihood estimation (CMLE) method. The standard approach to solve the CMLE problem is the Classification-Expectation-Maximization (CEM) algorithm [CG92]. Unfortunately, it is only guaranteed that the algorithm converges to some (possibly arbitrarily poor) stationary point of the objective function.

In this paper, we present two algorithms for a restricted version of the CMLE problem. That is, our algorithms approximate reasonable solutions to the CMLE problem which satisfy certain natural properties. Moreover, they compute solutions whose cost (i.e. complete-data log-likelihood values) are at most a factor $(1 + \varepsilon)$ worse than the cost of the solutions that we search for. Note the CMLE problem in its most general, i.e. unrestricted, form is not well defined and allows for trivial optimal solutions that can be thought of as degenerated solutions.

1 Preliminaries

Given set of observations, the objective of the CMLE problem is to find a Gaussian mixture model and a hard clustering with maximum complete-data likelihood. In this section, we will first describe and define this objective function. Then, we will present an alternating optimization scheme for this problem. However, the problem is not well-defined. Hence, we will restrict the problem to reasonable instances and solutions.

1.1 Complete-Data Log-Likelihood

Let $X \subset \mathbb{R}^d$ be a finite set of observations. Given a spherical Gaussian distribution $\mathcal{N}_d(\mu, \sigma)$, the *likelihood* that all $x \in X$ have been drawn according to $\mathcal{N}_d(\mu, \sigma)$ is given by

$$\prod_{x \in X} \mathcal{N}_d(x|\mu, \sigma) \ ,$$

assuming that the observations have been drawn independently at random.

Definition 1. *Given a finite set $X \subset \mathbb{R}^d$ and a spherical Gaussian distribution with mean $\mu \in \mathbb{R}^d$ and variance $\sigma^2 \in \mathbb{R}$, let*

$$\mathcal{L}_X(\mu, \sigma^2) \coloneqq -\ln\left(\prod_{x \in X} p(x|\mu, \sigma^2)\right) = \frac{|X|d}{2}\ln(2\pi\sigma_k^2) + \frac{1}{2\sigma_k^2}\sum_{x \in X}\|x - \mu_k\|^2 \ .$$

We denote the minimal value by $OPT(X, 1) = \min_{(\mu,\sigma^2)} \mathcal{L}_X(\mu, \sigma^2)$.

Now consider a Gaussian mixture model (GMM) given by parameters $\theta = \{(w_k, \mu_k, \sigma_k^2)\}_{k=1}^K$. Drawing an observation x_n according to a GMM corresponds to a two-step process:

1. Draw a component $z_n \in [K]$ with probability $p(z_n = k|\theta) = w_k$.
2. Draw an observation $x_n \in X$ according to $\mathcal{N}_d(\mu_{z_n}, \sigma_{z_n})$.

Note that the assignment $z_n \in [K]$ is a (latent) random variable in this two-step process. With the help of this random variable, we can compute the likelihood that observation $x \in X$ has been generated by the k-th component of the GMM, i.e.

$$p(x_n, z_n = k|\theta) = p(z_n = k|\theta) \cdot p(x_n|z_n = k, \theta) = w_k \cdot \mathcal{N}_d(x|\mu_k, \sigma_k) \ .$$

Since x_n and z_n completely describe the two-step process, the likelihood $p(x_n, z_n|\theta)$ is also called *complete-data likelihood*, while $p(x_n|\theta) = \sum_{z_n=1}^K p(x_n, z_n|\theta)$ is refered to as *(marginal) likelihood*.

Assume, we are given a set of observations $X = \{x_n\}_{n=1}^N$ and assignments $\{z_n\}_{n=1}^N$. Then, the likelihood that all observations have been drawn according to a GMM θ and that each x_n has been generated by the z_n-th component, is given by

$$\prod_{n=1}^N p(x_n, z_n|\theta) = \prod_{n=1}^N w_{z_n} \cdot \mathcal{N}_d(x_n|\mu_{z_n}, \sigma_{z_n}) \ , \tag{1}$$

assuming that the observations have been drawn inpendently at random. Note that the assignments $\{z_n\}_{n=1}^N$ define a partition $\mathcal{C} = \dot{\cup}_{k=1}^K C_k$ via $x_n \in C_k$ iff $z_n = k$. Hence, we can also rewrite Equation (1) as

$$\prod_{k=1}^K \prod_{x_n \in C_k} p(x_n, z_n = k|\theta) = \prod_{k=1}^K \prod_{x_n \in C_k} w_k \cdot \mathcal{N}_d(x_n|\mu_k, \sigma_k) \ .$$

By taking (negative) logarithm of this expression, we obtain

$$
\begin{aligned}
&-\log\left(\prod_{k=1}^{K}\prod_{x_n\in C_k} p(x_n, z_n = k|\theta)\right)\\
&=\sum_{k=1}^{K}\sum_{x_n\in C_k}\left(\ln(w_k)+\ln\left(\mathcal{N}_d(x_n|\mu_k,\Sigma_k)\right)\right)\\
&=\sum_{k=1}^{K}\mathcal{L}_{C_k}(\mu_k,\sigma_k^2)-\ln(w_k)\cdot|C_k|\ .
\end{aligned}
$$

Definition 2. *Given a finite set $X\subset\mathbb{R}^d$, a partition $\mathcal{C}=\{C_1,\dots,C_K\}$ of X, and a mixture of spherical Gaussians with parameters $\theta=\{(w_k,\mu_k,\sigma_k^2)\}_{k=1}^K$, we call*

$$\mathcal{L}_X(\theta,\mathcal{C}) \coloneqq \sum_{k=1}^{K}\mathcal{L}_{C_k}(\mu_k,\sigma_k^2)-\ln(w_k)\cdot|C_k|$$

the complete-data negative log-likelihood.

Note that a solution maximizing the complete-data likelihood also minimizes the complete-data negative log-likelihood, and vice versa. Therefore, we define the *complete-cata maximum likelihood estimation* (CMLE) problem as follows.

Problem 3 (CMLE)**.** *Given a finite set $X\subset\mathbb{R}^d$ and an integer $K\in\mathbb{N}$, find a partition $\mathcal{C}=\{C_1,\dots,C_K\}$ of X and a mixture of spherical Gaussians with parameters $\theta=\{(w_k,\mu_k,\sigma_k^2)\}_{k=1}^K$ minimizing $\mathcal{L}_X(\theta,\mathcal{C})$. We denote the minimal value by $OPT(X,K)$.*

For a fixed model θ, we let $\mathcal{L}_X(\theta)=\min_{\mathcal{C}}\mathcal{L}_X(\theta,\mathcal{C})$. Analogously, for a fixed clustering $\mathcal{C}$, we let $\mathcal{L}_X(\mathcal{C})=\min_{\theta}\mathcal{L}_X(\theta,\mathcal{C})$.

Definition 4. *Given parameters (w_k,μ_k,σ_k^2) and a cluster $C_k\subseteq X$, we let*

$$\mathcal{L}_x(w_k,\mu_k,\sigma_k^2) \coloneqq \frac{d}{2}\ln(2\pi\sigma_k^2)+\frac{1}{2\sigma_k^2}\|x-\mu_k\|^2-\ln(w_k),$$

and

$$\mathcal{L}_{C_k}(w_k,\mu_k,\sigma_k^2) \coloneqq \sum_{x\in C_k}\mathcal{L}_x(w_k,\mu_k,\sigma_k^2)\ .$$

Remark 5. *For all partitions $\mathcal{C}=\{C_1,\dots,C_K\}$, we have*

$$\mathcal{L}_X(\mathcal{C})=\sum_{k=1}^{K}OPT(C_k,1)-\ln\left(\frac{|C_k|}{|X|}\right)\cdot|C_k|\ .$$

For all $\theta=\{(w_1,\mu_1,\sigma_1^2),\dots,(w_K,\mu_K,\sigma_K^2)\}$, we have

$$\mathcal{L}_X(\theta)=\sum_{n=1}^{N}\mathrm{argmin}_{k\in[K]}\{\mathcal{L}_x(w_k,\mu_k,\sigma_k^2)\}\ .$$

1.2 Alternating Optimization Scheme (CEM algorithm)

An *alternating optimization algorithm* for this problem is given by the following first order optimality conditions. Fixing the partition $\mathcal{C} = \{C_k\}_{k=1}^K$, the optimal mixture of spherical Gaussians is given by $\theta = \{(w_k, \mu_k, \sigma_k^2)\}_{k=1}^K$ with

$$w_k = \frac{|C_k|}{|X|} \ , \qquad \mu_k = \frac{1}{|C_k|} \sum_{x_n \in C_k} x_n \ , \qquad \sigma_k^2 = \frac{1}{d|C_k|} \sum_{x_n \in C_k} \|x_n - \mu_k\|^2 \ .$$

Fixing the Gaussian mixture model $\theta = \{(w_k, \mu_k, \sigma_k^2)\}_{k=1}^K$, the optimal partition $\mathcal{C} = \{C_k\}_{k=1}^K$ is given by assigning each point to its most likely component, i.e.

$$x_n \in C_k \Leftrightarrow k = \text{argmax}_{l \in [K]}\, p(z_n = l | x_n, \theta) \ ,$$

where

$$p(z_n = k | x_n, \theta) = \frac{w_k \mathcal{N}(x_n | \mu_k, \sigma_k^2)}{\sum_{l=1}^K w_l \mathcal{N}(x_n | \mu_l, \sigma_l^2)} \ ,$$

which is the *posterior probability* that x_n has been generated by the k-th component of the given mixture.

If we repeatedly compute these update formulas, the solution converges to a local extremum or a saddlepoint of the likelihood function.

A proof of the correctenss of these update formulas (which we omit here) uses the following lemma.

Lemma 6. *Let $X \subset \mathbb{R}^d$ be a finite set. Define*

$$\mu(X) = \frac{1}{|X|} \sum_{x \in X} x \ .$$

Then, for all $y \in \mathbb{R}^d$

$$\sum_{x \in X} \|x - y\|^2 = \sum_{x \in X} \|x - \mu(X)\|^2 + |X| \cdot \|y - \mu(X)\|^2 \ .$$

In particular, $\mu(X) = \text{argmin}_{y \in \mathbb{R}^d} \sum_{x \in X} \|x - y\|^2$.

Note that an optimal CMLE solution is not changed by this algorithm. Hence, an optimal CMLE solution is completely defined by a partition or a Gaussian mixture model. Similarly, if we refer to a partition or a Gaussian mixture as a CMLE solution we assume that the missing parameters are as defined by the update formulas given above, respectively.

1.3 Well-Defined Instances

Unfortunately, the CMLE problem is not well defined in this form. For example, you could choose $C_1 = \{x\}$ and $\mu_1 = x$ for some $x \in X$. Then, as $\sigma_1 \to 0$ we get that $\mathcal{L}_K(X) \to -\infty$. Consequently, we impose the following restrictions on instances.

Definition 7. *We call $X = \dot{\bigcup}_{k=1}^K C_k$ a* well-defined partition *if*

1. *for all $k \in [K]$: $|C_k| \geq 2$.*

We call X itself a well-defined instance *if*

2. $\forall x, y \in X, x \neq y: \ \|x-y\|^2 \geq \frac{4d}{\pi}$.

We denote $X = \dot{\bigcup}_{k=1}^{K} C_k$ *as a* well-defined solution *if* X *is a well-defined instance and* $\{C_k\}_{k=1}^{K}$ *is a well-defined partition.*

In the following, we prove that, with these restrictions, the CMLE problem is well defined. That is, the minimum in Problem 3 is well defined ($\mathcal{L}_K(X) > -\infty$). Moreover, we will see (Lemma 9) that for the optimal solution we have $\sigma_k^2 \geq \frac{1}{2\pi}$ or

$$2\pi\sigma_k^2 \geq 1 \quad \text{for } k \in [K]. \tag{2}$$

First of all, note that the sum of squared distances between the points in X and the mean $\mu(X)$ can be rewritten using pairwise distances (which are lower bounded in Restriction 2).

Lemma 8. *Let* $X \subset \mathbb{R}^d$ *be a finite set and* $\mu(X) := \frac{1}{|X|}\sum_{x\in X}$ *its mean, then*

$$\sum_{x\in X}\|x-\mu(X)\|^2 = \frac{1}{2|X|}\sum_{x\in X}\sum_{y\in X}\|x-y\|^2.$$

Proof.

$$\begin{aligned}
\sum_{x\in X}\sum_{y\in X}\|x-y\|^2 &= \sum_{x\in X}\sum_{y\in X}\langle x-y, x-y\rangle \\
&= \sum_{x\in X}\sum_{y\in X}(\langle x,x\rangle + \langle y,y\rangle - 2\langle x,y\rangle \\
&= 2|X|\sum_{x\in X}\langle x,x\rangle - 2\sum_{x\in X}\sum_{y\in X}\langle x,y\rangle \\
&= 2|X|\sum_{x\in X}\langle x,x\rangle - 2|X|\sum_{x\in X}\langle x,\mu(X)\rangle \\
&= 2|X|\sum_{x\in X}\langle x, x-\mu(X)\rangle \\
&= 2|X|\sum_{x\in X}\langle x-\mu(X), x-\mu(X)\rangle \qquad (\text{using } |X|\textstyle\sum_{x\in X}\langle \mu(X), x-\mu(X)\rangle = 0) \\
&= 2|X|\sum_{x\in X}\|x-\mu(X)\|^2.
\end{aligned}$$

□

Now using the restriction on the minimum pairwise difference between points (Restriction 2) and on the minimum number of points (Restriction 1) in a cluster, we can lower bound the variance of each cluster. This directly yields Equation (2) and our claim that the problem is well-defined under the restrictions given in Definition 7.

Lemma 9. *Let* Y *be a subset of a set* X *that satisfies Restriction 2 from Definition 7 and that contains at least two different elements. Then,*

$$\sigma(Y)^2 = \frac{1}{|Y|d}\sum_{y\in Y}\|y-\mu(Y)\|^2 \geq \frac{1}{2\pi}\ .$$

Proof.

$$\begin{aligned}
\sigma(Y)^2 &= \frac{1}{|Y|d}\sum_{y\in Y}\|y-\mu(Y)\|^2 \\
&= \frac{1}{2|Y|^2 d}\sum_{x\in Y}\sum_{y\in Y}\|x-y\|^2 && \text{(using Lemma 8)} \\
&\geq \frac{1}{2|Y|^2 d}\binom{|Y|}{2}\min_{x,y\in Y, x\neq y}\|x-y\|^2 \\
&\geq \frac{1}{8d}\min_{x,y\in Y, x\neq y}\|x-y\|^2 \\
&\geq \frac{1}{2\pi} && \text{(using Restriction 2)}
\end{aligned}$$

□

Throughout the rest of this paper, we will restrict the search space of CMLE to well-defined solutions. In particular, we only consider the optimal solution among all well-defined solutions.

1.4 Well-Balanced Instances

A central idea behind the algorithms that we present in this paper is that we do not allow somewhat *degenerate* instances. This means that we can find a function f in the number of clusters that can be used to lower bound the number of points in a cluster and a function g that can be used to lower bound the costs $OPT(C_k, 1)$ of optimal clusters C_k.

Definition 10 (well-balanced)**.** *Let $f, g : \mathbb{N} \to \mathbb{R}$. We denote a partition $X = \dot{\bigcup}_{k=1}^{K} C_k$ as* f-balanced *if for all $k \in [K]$*

$$|C_k| \geq \frac{|X|}{f(K)} \ .$$

Furthermore, we denote the partition as an (f, g)-balanced *CMLE solution if it is f-balanced and additionally for all $k \in [K]$*

$$OPT(C_k, 1) \geq \frac{1}{g(K)} \cdot \sum_{k=1}^{K} OPT(C_k, 1) \ .$$

Definition 11. *Given a finite set $X \subset \mathbb{R}^d$ and $K \in \mathbb{N}$, we let*

$$OPT_{diam}(X, K) = \min_{\substack{\{C_1,\dots,C_K\},\\ \dot{\cup}_{k=1}^{K} C_k = X}} \max_{k\in[K]} \max_{x,y\in C_k} \|x-y\| \ .$$

Lemma 12 (From f-balanced to (f, g)-balanced)**.** *An f-balanced solution $X = \dot{\bigcup}_{k=1}^{K} C_k$ is also an $(f, \Gamma \cdot f)$-balanced CMLE solution, where $\Gamma \leq 2 \cdot \ln(32\pi \cdot OPT_{diam}(X, K)) + \ln(K) + 1$.*

Proof.

$$\begin{aligned}
OPT(C_k, 1) &\geq \frac{|C_k|d}{2} \geq \frac{1}{f(K)}\frac{|X|d}{2} && \text{(due to Lemma 20 and } f \text{ balanced)} \\
&\geq \frac{1}{f(K)\cdot\Gamma}\mathcal{L}_K(X) && \text{(due to Lem. 21)} \\
&\geq \frac{1}{f(K)\cdot\Gamma}\sum_{k=1}^{K} OPT(C_k, 1) \ .
\end{aligned}$$

□

2 Main Results (Theorems 13 and 15)

Theorem 13. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$ and $\delta, \varepsilon \in [0,1]$. If X has an (f,g)-balanced optimal CMLE solution, then there exists an algorithm which computes a mixture of K spherical Gaussians $\theta = \{(w_k, \mu_k, \sigma_k^2)\}_{k=1}^K$, such that*

$$Pr\left[\mathcal{L}_X(\theta) \leq (1+\varepsilon)OPT(X,K)\right] \geq 1-\delta \ .$$

The runtime of the algorithm is bounded by

$$|X| \cdot K \cdot \log(\Gamma) \cdot \log(g(K)) \cdot 2^{\tilde{\mathcal{O}}\left(\frac{f(K)}{\varepsilon\delta}\right)}$$

where $\Gamma \leq 2 \cdot \ln\left(32\pi \cdot OPT_{diam}(X,K)\right) + \ln(K) + 1$.

Corollary 14. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$ and $\delta, \varepsilon \in [0,1]$. If X has an f-balanced optimal CMLE solution, then there exists an algorithm which computes a mixture of K spherical Gaussians θ, such that*

$$Pr\left[\mathcal{L}_X(\theta) \leq (1+\varepsilon)OPT(X,K)\right] \geq 1-\delta \ .$$

The runtime of the algorithm is bounded by

$$|X| \cdot K \cdot \log(\Gamma)^2 \cdot 2^{\tilde{\mathcal{O}}\left(\frac{f(K)}{\varepsilon\delta}\right)}$$

where $\Gamma \leq 2 \cdot \ln\left(32\pi \cdot OPT_{diam}(X,K)\right) + \ln(K) + 1$.

Theorem 15. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$, and $\delta, \varepsilon > 0$. Let $\mathcal{C} = \dot{\bigcup}_{k=1}^K C_k$ be a well-defined solution for the CMLE problem. There is an algorithm that computes a mixture of K spherical Gaussians θ, such that*

$$\Pr\left[\mathcal{L}_X(\theta) \leq (1+\varepsilon)\mathcal{L}_X(\mathcal{C})\right] \geq 1-\delta \ .$$

The running time of the algorithm is bounded by

$$|X|\, d\, \log\left(\frac{1}{\delta}\right) 2^{\mathcal{O}\left(\frac{K}{\varepsilon}\cdot\log\left(\frac{K}{\varepsilon^2}\right)\right)} \left(\log(\log(\Delta^2))+1\right)^K \left(\log(f(K))\right)^K \ ,$$

where $\Delta^2 = \max_{x,y\in X}\{\|x-y\|^2\}$.

3 Proof of Theorem 13

In the following we prove Theorem 13.

- In Section 3.1 we show that, if the parameters of a CMLE solution are sufficently close to those of an optimal CMLE solution, then its complete-data log-likelihood is close to that of the optimal CMLE solution. In Sections 3.2 and 3.3 we then show how to obtain such parameter estimates.
- In Section 3.2 we deal with the problem of estimating the means. We use the superset sampling technique introduced by [IKI94] to compute a set of candidate means which contains a good candidate, i.e. a good estimation to the mean parameters of an optimal solution.

- In Section 3.3 we use a grid search to obtain estimates of the weights and variances. The core idea is to simply test all solutions lying on a specific grid in the search space. By choosing a grid that is dense enough, we ensure that there are solutions on the grid which are sufficiently close to the parameters that we search for.

3.1 Estimate the Costs of Parameter Estimates

For an optimal (f, g)-balanced CMLE solutions, we can estimate the parameters of the the respective optimal Gaussian mixture model and the likelihood of the optimal clusters. We can show that the CMLE solution determined by these parameter estimates yields an approximation with respect to the complete data log-likelihood.

Theorem 16. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$ and $\varepsilon > 0$. Assume X has an f-balanced optimal CMLE solution $X = \dot{\bigcup}_{k=1}^{K} C_k$ and let $(\tilde{\mu}_1, \ldots, \tilde{\mu}_K)$ such that for all $k \in [K]$*

$$\|\tilde{\mu}_k - \mu(C_k)\|^2 \leq \frac{\varepsilon}{|C_k|} \sum_{x \in C_k} \|x - \mu(C_k)\|^2 \ .$$

Let $(n_1, \ldots, n_K)$, such that for all $k \in [K]$

$$|C_k| \leq n_k \leq (1+\varepsilon)|C_k| \ . \tag{3}$$

and $\vec{\tilde{\sigma}} = (\tilde{\sigma}_1^2, \ldots, \tilde{\sigma}_K^2) \in \mathbb{R}^K$, such that for all $k \in [K]$ it holds

$$\tilde{\sigma}_k^2 \geq \sigma_k^2 \tag{4}$$

and

$$\ln(\tilde{\sigma}_k^2) - \ln(\sigma_k^2) \leq \left((1+\varepsilon)^2 - 1\right) \frac{2}{|C_k| d} OPT(C_k, 1) \ . \tag{5}$$

Define $\tilde{\theta} = \{(\tilde{w}_k, \tilde{\mu}_k, \tilde{\sigma}_k^2)\}_{k=1,\ldots,K}$, where $\tilde{w}_k = \frac{n_k}{\sum_{l=1}^{K} n_l}$. Then,

$$\mathcal{L}_X(\tilde{\theta}) \leq (1+\varepsilon)^4 OPT(X, K).$$

Proof. Using that $|C_l| \leq n_l \leq (1+\varepsilon)|C_l|$ for all $l = 1, \ldots, K$, we obtain $\tilde{w}_k \geq \frac{1}{(1+\varepsilon)} \cdot \frac{|C_k|}{|X|}$. Hence,

$$\begin{aligned}
-\ln(\tilde{w}_k) \cdot |C_k| &\leq -\ln\left(\frac{1}{(1+\varepsilon)} \cdot \frac{|C_k|}{|X|}\right) |C_k| && \text{(by Equation (3))} \\
&\leq \ln(1+\varepsilon)|C_k| - \ln\left(\frac{|C_k|}{|X|}\right) \cdot |C_k| \\
&\leq \varepsilon |C_k| - \ln\left(\frac{|C_k|}{|X|}\right) \cdot |C_k| && \text{(since } \ln(1+\varepsilon) \leq \varepsilon\text{)} \\
&\leq \frac{2\varepsilon}{d} OPT(C_k, 1) - \ln\left(\frac{|C_k|}{|X|}\right) \cdot |C_k| && \text{(since } OPT(C_k, 1) \geq \tfrac{|C_k| \cdot d}{2}\text{)}
\end{aligned}$$

Furthermore, observe that

$$\begin{aligned}
\mathcal{L}_{C_k}(\tilde{\mu}_k, \tilde{\sigma}_k) &= \frac{|C_k|d}{2}\ln(2\pi\tilde{\sigma}_k^2) + \frac{1}{2\tilde{\sigma}_k^2}\sum_{x\in C_k}\|x-\tilde{\mu}_k\|^2 \\
&\overset{(4)}{\leq} \frac{|C_k|d}{2}\ln(2\pi\tilde{\sigma}_k^2) + \frac{1}{2\sigma_k^2}\sum_{x\in C_k}\|x-\tilde{\mu}_k\|^2 \\
&\leq \frac{|C_k|d}{2}\ln(2\pi\tilde{\sigma}_k^2) + \frac{1}{2\sigma_k^2}(1+\varepsilon)\sum_{x\in C_k}\|x-\mu_k\|^2 \quad \text{(By Lemma 6 and property of } \tilde{\mu}_k) \\
&= \frac{|C_k|d}{2}\ln(2\pi\tilde{\sigma}_k^2) + (1+\varepsilon)\frac{|C_k|d}{2} \qquad \text{(By def. of } \mu_k) \\
&= \frac{|C_k|d}{2}(\ln(2\pi) + \ln(\tilde{\sigma}_k^2)) + (1+\varepsilon)\frac{|C_k|d}{2} \\
&\overset{(5)}{=} \frac{|C_k|d}{2}\left(\ln(2\pi) + \left((1+\varepsilon)^2-1\right)\frac{2}{|C_k|d}OPT(C_k,1) + \ln(\sigma_k^2)\right) + (1+\varepsilon)\frac{|C_k|d}{2} \\
&= \frac{|C_k|d}{2}\ln(2\pi\sigma_k^2) + (1+\varepsilon)\frac{|C_k|d}{2} + \left((1+\varepsilon)^2-1\right)OPT(C_k,1) \\
&\leq (1+\varepsilon)OPT(C_k,1) + \left((1+\varepsilon)^2-1\right)OPT(C_k,1) \\
&\leq \left((1+\varepsilon)^2+\varepsilon\right)OPT(C_k,1) \\
&\leq (1+\varepsilon)^3 OPT(C_k,1)
\end{aligned}$$

Overall, we have

$$\begin{aligned}
\mathcal{L}_X(\tilde{\theta}) &= \sum_{k=1}^{K}\mathcal{L}_{C_k}(\mu_k,\sigma_k^2) - \ln(w_k)\cdot|C_k| \\
&\leq \sum_{k=1}^{K}(1+\varepsilon)^3 OPT(C_k,1) + \frac{2\varepsilon}{d}OPT(C_k,1) - \ln\left(\frac{|C_k|}{|X|}\right)\cdot|C_k| \\
&= \sum_{k=1}^{K}\left((1+\varepsilon)^3+\frac{2\varepsilon}{d}\right)OPT(C_k,1) - \ln\left(\frac{|C_k|}{|X|}\right)\cdot|C_k| \\
&\leq \left((1+\varepsilon)^3+\frac{2\varepsilon}{d}\right)\sum_{k=1}^{K}OPT(C_k,1) - \ln\left(\frac{|C_k|}{|X|}\right)\cdot|C_k| \\
&= \left((1+\varepsilon)^3+\frac{2\varepsilon}{d}\right)OPT(X,K) \\
&\leq (1+\varepsilon)^4 OPT(X,K)
\end{aligned}$$

□

3.2 Generate Candidate Means by Sampling

We reuse the following well-known lemma on superset sampling.

Lemma 17 (superset-sampling)**.** *Let $X \subset \mathbb{R}^d$ be a finite set, $\alpha < 1$ and $X' \subset X$ with $|X'| \geq \alpha|X|$. Let $S \subseteq X$ be a uniform sample multiset of size at least $\frac{2}{\alpha\varepsilon\delta}$. Then with probability at*

least $\frac{1-\delta}{5}$ *there is a subset* $S' \subseteq S$ *with* $|S'| = \frac{1}{\varepsilon\delta}$ *such that*

$$\|\mu(S') - \mu(X')\|^2 \leq \frac{\varepsilon}{|X'|} \sum_{x \in X'} \|x - \mu(X')\|^2.$$

If we plug our notion of f-balanced solutions into this lemma, then we receive an algorithm that samples good approximative means.

Theorem 18 (sampling means)**.** *For a finite set* $X \subset \mathbb{R}^d$*,* $K \in \mathbb{N}$ *and* $\varepsilon, \delta > 0$*, if* $X = \dot{\bigcup}_{k=1}^{K} C_k$ *is an* f*-balanced partition, then there is an algorithm that computes a set of* $\log(1/\delta) \cdot 2^{\frac{K}{\varepsilon\delta} \cdot \log\left(\frac{f(K)}{\varepsilon\delta}\right)}$ K*-tuples of points from* $\mathbb{R}^d$*, such that with probability* $1 - \delta$ *for one of these tuples it holds that for all* $k \in [K]$

$$\|\mu_k - \mu(C_k)\|^2 \leq \frac{\varepsilon}{|C_k|} \sum_{x \in C_k} \|x - \mu(C_k)\|^2 \ .$$

The runtime of the algorithm is bounded by $\log(1/\delta) \cdot K \cdot \left(|X| + 2^{\frac{K}{\varepsilon\delta} \cdot \log\left(\frac{f(K)}{\varepsilon\delta}\right)}\right)$.

Proof. Consider the following algorithm, which computes a candidate set of tuples of means.

Algorithm 1: Approx-Means(X, K)

Input: $X \subset \mathbb{R}^d$: input points
$K \in \mathbb{N}$: number of clusters
Output: set of candidate tuples of means
$P \leftarrow \emptyset$;
for $k = 1, \ldots, K$ **do**
  sample a multiset S of size $\frac{1}{\alpha\varepsilon\delta}$ from X;
  $T \leftarrow \left\{\mu(S') | S' \subset S, |S'| = \lceil \frac{1}{\varepsilon\delta} \rceil\right\}$;
  $P \leftarrow P \times T$;
end
return P;

Using Lemma 17 with $\alpha = \frac{1}{f(K)}$, we know that the output of a single run of Approx-Means contains a tuple with the desired property with probability $\left(\frac{1-\delta}{5}\right)^K$.

We know that

$$|T| \leq \left(\frac{1}{\alpha\varepsilon\delta}\right)^{\frac{1}{\varepsilon\delta}},$$

thus

$$|P| = |T|^K \leq 2^{\frac{K}{\varepsilon\delta} \cdot \log\left(\frac{f(K)}{\varepsilon\delta}\right)}.$$

The runtime is bounded by

$$K \cdot |X| + \sum_{k=1}^{K} |T|^k \leq K \left(|X| + 2^{\frac{K}{\varepsilon\delta} \cdot \log\left(\frac{f(K)}{\varepsilon\delta}\right)}\right).$$

By executing Approx-Means $\log(1/\delta)$ times we receive the desired success probability.

□

3.3 Generate Candidate Cluster Sizes and Variances by Using Grids

So far, we have formulated an algorithm that gives us good means. In the following, we will use the gridding technique to determine a set of candidates for the the cluster sizes and variances. First of all, we generate a set of cluster sizes that contains good approximations of the cluster sizes of any f-balanced solutions. Then, we approximate the negative log-likelihood of optimal CMLE clusters, i.e. $\sum_{k=1}^{K} OPT(C_k, 1)$ where the C_k are the optimal CMLE clusters. Then, we present how to construct a candidate set of variances that contains good estimates of the variances of any (f, g)-balanced optimal CMLE solution.

3.3.1 Grid Search for Cluster Sizes

Theorem 19. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$ and let $X = \dot{\bigcup}_{k=1}^{K} C_k$ be an f-balanced partition. Then there exists an algorithm that outputs a set $S \subseteq \mathbb{N}^K$, $|S| = \left(\frac{\log(f(K))}{\log(1+\varepsilon)}\right)^K$, that contains a tuple $(n_1, \ldots, n_K) \in S$ such that*

$$|C_k| \leq n_k \leq (1+\varepsilon)|C_k|. \tag{6}$$

for all $k \in [K]$.

Proof. Since we assume a f-balanced solution, we know that for all $k \in [K]$

$$\frac{|X|}{f(K)} \leq |C_k| \leq |X|.$$

Thus, there exist a value $i^* \in \{1, \ldots, \lceil \log_{1+\varepsilon}(f(K)) \rceil\}$ such that

$$(1+\varepsilon)^{i^*-1} \frac{|X|}{f(K)} \leq |C_k| \leq (1+\varepsilon)^{i^*} \frac{|X|}{f(K)}.$$

Thus, we receive $\lceil \log_{1+\varepsilon}(f(K)) \rceil$ many values for each cluster size n_k. The algorithm outputs all possible combinations of these values. □

3.3.2 Bounds on the Log-Likelihood of optimal CMLE clusters

Lemma 9 provides us with a lower bound on the negative log-likelihood of a cluster.

Corollary 20 (Lower Bound on the Optimal Log-Likelihood)**.** *Let $X = \dot{\bigcup}_{k=1}^{K} C_k$ be an optimal CMLE solution. Then, $OPT(C_k, 1) \geq \frac{|C_k|d}{2}$.*

The next step is to find an upper bound on the optimal complete-data likelihood value. We use Gonzales algorithm to compute a value that gives us a tighter bound than just the maximum spread (over the dimensions of the vectors in the data set).

Lemma 21 (Upper Bound on the Optimal Complete-Data Log-Likelihood)**.** *Let $X \subset \mathbb{R}^d$ and $K \in \mathbb{N}$. A Value Γ can be computed in time $\mathcal{O}(K \cdot d \cdot |X|)$ such that the complete-data likelihood of an optimal CMLE solution can be bounded by*

$$OPT(X, K) \leq \frac{|X|d}{2} \cdot \Gamma$$

and $\Gamma = \ln(2\pi s^2) + 1 + \ln(K)$ for some $s \leq 4 \cdot OPT_{diam}(X)$.

Proof. Run Gonzales algorithm. The output is a set of K points $p_1, \ldots, p_K \in X$. Compute the point z with maximum distance to its closest point in $\{p_1, \ldots, p_K\}$ and set $s := \min_{k=1,\ldots,K} \|z - p_k\|$. Consider the solution where the p_k are the centers. Partition the points into point sets $\mathcal{C} = \{C_1, \ldots, C_K\}$, with $\|x - p_k\| = \min_{i=1,\ldots,K} \|x - p_i\|$ for all $x \in C_k$. Notice that the distances between any point and its center is at most s. Thus, when computing the optimal variance in each cluster, it is at most s^2. Then, for $\theta = \left\{\left(\frac{1}{K}, p_k, \sigma(X_k, p_k)\right)\right\}_{k=1}^{K}$ we have

$$\begin{aligned} OPT(X,K) \leq \mathcal{L}_X(\theta, \mathcal{C}) &= \sum_{k=1}^{K} \frac{|C_k|d}{2} \ln(2\pi\sigma(C_k, p_k)^2) + \frac{|C_k|d}{2} - \ln(w_k) \cdot |C_k| \\ &\leq \left(\sum_{k=1}^{K} \frac{|C_k|d}{2} \ln(2\pi s^2) + \frac{|C_k|d}{2} \right) - \ln\left(\frac{1}{K}\right) \cdot |X| \\ &= \frac{|X|d}{2} \ln(2\pi s^2) + \frac{|X|d}{2} + \ln(K) \cdot |X| \\ &\leq \frac{|X|d}{2} \left(\ln(2\pi s^2) + 1 + \ln(K)\right) \end{aligned}$$

□

Given two bounds, we can find a constant factor approximation of the the sum of the negative log-likelihoods of optimal CMLE clusters, i.e. $\sum_{k=1}^{K} OPT(C_k, 1)$, using a grid search.

Lemma 22 (Estimating the Optimal Log-Likelihood)**.** *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$, and $\varepsilon > 0$. Let $X = \dot{\cup}_{k=1}^{K} C_k$ be an optimal CMLE solution. Then, there exists a set of $\log(3\Gamma/d)/\log(1+\varepsilon)$ many values which contains a value $\mathcal{N}_{est}$ with*

$$\frac{1}{1+\varepsilon} \mathcal{N}_{est} \leq \sum_{k=1}^{K} OPT(C_k, 1) \leq \mathcal{N}_{est} \ .$$

Proof. Combining Corollary 20 and Lemma 21, we know that

$$\frac{|X|d}{2} \leq \sum_{k=1}^{K} OPT(C_k, 1) \leq OPT(X, K) \leq \frac{|X|d}{2} \Gamma.$$

Thus, there exist a value $i^* \in \{1, \ldots, \lceil \log_{1+\varepsilon}(\Gamma) \rceil\}$ such that

$$(1+\varepsilon)^{i^*-1} \frac{|X|d}{2} \leq \sum_{k=1}^{K} OPT(C_k, 1) \leq (1+\varepsilon)^{i^*} \frac{|X|d}{2}.$$

The algorithm outputs all $\lceil \log_{1+\varepsilon}(\Gamma) \rceil$ values. □

Given this approximation of the sum of the negative log-likelihoods, we will be able to find an approximation of the negative log-likelihoods of a single cluster as we will see in the next section.

3.3.3 Grid Search for Variances

Given the approximations of the size of the clusters and their negative log-likelihod, we are now able to find estimates of the variances.

Theorem 23. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$ and $\varepsilon > 0$. Assume X has an (f,g)-balanced CMLE solution $X = \dot{\bigcup}_{k=1}^{K} C_k$. Let additionally $\mathcal{N}_{est} \in \mathbb{R}$, with*

$$\frac{1}{1+\varepsilon}\mathcal{N}_{est} \leq \sum_{k=1}^{K} OPT(C_k, 1) \leq \mathcal{N}_{est}, \tag{7}$$

and $(n_1, \ldots, n_K)$, such that for all $k \in [K]$

$$|C_k| \leq n_k \leq (1+\varepsilon)|C_k|. \tag{8}$$

Then there exists an algorithm that computes a set of size $K \cdot \frac{\log(g(K))}{\log(1+\varepsilon)}$, that contains a tuple $(\tilde{\sigma}_1^2, \ldots, \tilde{\sigma}_K^2)$, such that for all $k \in [K]$ it holds

$$\tilde{\sigma}_k^2 \geq \sigma_k^2 \tag{9}$$

and

$$\ln(\tilde{\sigma}_k^2) - \ln(\sigma_k^2) \leq \left((1+\varepsilon)^2 - 1\right) \frac{2}{|C_k|d} OPT(C_k, 1)\ . \tag{10}$$

Proof. Observe that

$$\frac{1}{g(K)(1+\varepsilon)}\mathcal{N}_{est} \leq \frac{1}{g(K)} \sum_{k=1}^{K} OPT(C_k, 1) \overset{\text{Def. } 10}{\leq} OPT(C_k, 1) \leq \sum_{k=1}^{K} OPT(C_k, 1) \leq \mathcal{N}_{est}.$$

Thus, there exists a value $j^* \in \left\{\lceil -\log_{1+\varepsilon}(g(K))\rceil, \ldots, 0\right\}$ which satisfies

$$(1+\varepsilon)^{j^*-1}\mathcal{N}_{est} \leq OPT(C_k, 1) \leq (1+\varepsilon)^{j^*}\mathcal{N}_{est}\ .$$

Denote the upper bound by $\hat{\mathcal{N}} \coloneqq (1+\varepsilon)^{j^*}\mathcal{N}_{est}$ and set $\tilde{\sigma}_k^2 \coloneqq \exp\left(\frac{2(1+\varepsilon)}{n_k d}\hat{\mathcal{N}} - \ln(2\pi) - 1\right)$.

Notice that

$$\begin{aligned} OPT(C_k, 1) = \mathcal{L}_{C_k}(\mu_k, \sigma_k^2) &= \frac{|C_k|d}{2}\left(\ln(2\pi\sigma_k^2 + 1)\right) \\ \Leftrightarrow \ln(\sigma_k^2) &= \frac{2}{|C_k|d} OPT(C_k, 1) - \ln(2\pi) - 1 \end{aligned}$$

Thus,

$$\ln(\tilde{\sigma}_k^2) = \frac{2(1+\varepsilon)}{n_k d}\hat{\mathcal{N}} - \ln(2\pi) - 1 \geq \frac{2}{|C_k|d} OPT(C_k, 1) - \ln(2\pi) - 1 = \ln(\sigma_k^2)$$

and

$$\begin{aligned} \ln(\tilde{\sigma}_k^2) - \ln(\sigma_k^2) &= \frac{2(1+\varepsilon)}{n_k d}\hat{\mathcal{N}} - \frac{2}{|C_k|d} OPT(C_k, 1) \\ &\leq \frac{2(1+\varepsilon)^2}{|C_k|d} OPT(C_k, 1) - \frac{2}{|C_k|d} OPT(C_k, 1) \\ &= \left((1+\varepsilon)^2 - 1\right) \frac{2}{|C_k|d} OPT(C_k, 1) \end{aligned}$$

□

4 Proof of Theorem 15

In the following we present the proof of Theorem 15.

- In Section 4.1 we show how to estimate the variances and the cluster sizes of a well-defined CMLE solution via gridding. The idea behind a grid search is simply to test all solutions lying on a grid in the search space. By choosing a grid that is dense enough, we ensure that there are solutions on the grid which are sufficiently close to the parameters that we search for.
- In Section 4.2, we show how one can find good estimates of the means when given good estimates of the weights and covariances. To this end, we adapt the sample-and-prune technique presented in [ABS10].

4.1 Generate Candidates for Variances and Weights

Lemma 24. *Let $X \subset \mathbb{R}^d$, and $\{C_k\}_{k=1}^K$ be a well-defined CMLE solution for X, with corresponding variances $\{\sigma_k^2\}_{k=1}^K$. Then, there exists an algorithm which outputs a set of at most $\left(\frac{\log(\log(\Delta^2))+1}{\log(1+\varepsilon)}\right)^K$ tuples of variances, which contains a tuple $(\tilde{\sigma}_k^2)_{k=1}^K$, such that*

$$\forall k \in [K] : \sigma_k^2 \leq \tilde{\sigma}_k^2 \leq (\sigma_k^2)^{(1+\varepsilon)} \ ,$$

where $\Delta^2 = \max_{x,y \in X}\{\|x-y\|^2\}$.

Proof. We know that optimal variances σ_k^2 of a well-defined solution are bounded from below by

$$\forall k \in [K] : \frac{1}{2\pi} \leq \sigma_k^2.$$

Furthermore, we know that these are also bounded from above by

$$\forall k \in [K] : \sigma_k^2 = \frac{1}{|C_k|d} \sum_{x \in C_k} \|x - \mu(C_k)\|^2 \leq \frac{1}{|C_k|d} \sum_{x \in C_k} \Delta^2 \leq \Delta^2 \ .$$

Because $1/(2\pi) \leq \sigma_k^2 \leq \Delta^2$, there exists a value

$$k^* \in \{1, \ldots, \log_{1+\varepsilon}(-\log_{1/(2\pi)}(\Delta^2))\}$$

such that

$$(1/(2\pi))^{(1+\varepsilon)^{k^*-1}} \leq \sigma_i^2 \leq (1/(2\pi))^{(1+\varepsilon)^{k^*}} \ .$$

Thus, we receive $\left\lceil \frac{\log(\log(\Delta^2)) - \log(\log(2\pi))}{\log(1+\varepsilon)} \right\rceil$ many values for each variance. The algorithm outputs all possible combinations of these values. □

The following result is the same as in Section 3.3.

Theorem 25. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$ and let $\mathcal{C} = \dot{\bigcup}_{k=1}^K C_k$ be an f-balanced partition. Then there exists an algorithm that outputs a set $S \subseteq \mathbb{N}^K$, $|S| = \left(\frac{\log(f(K))}{\log(1+\varepsilon)}\right)^K$, that contains $\{n_1, \ldots, n_K\} \subset S$ such that*

$$|C_k| \leq n_k \leq (1+\varepsilon)|C_k|. \tag{11}$$

for all $k \in [K]$.

Algorithm 2: APPROX-MEANS$(R, l, \mathcal{M}_{k-l}, \Sigma)$

Input:

$R \subset X \subset \mathbb{R}^d$: set of remaining input points

$l \in \mathbb{N}$: number of means yet to be found

$\vec{\mu} = (\mu_1, \ldots, \mu_j)$: tuple of $j \leq k - l$ candidate means

$(\tilde{\sigma}_1^2, \ldots, \tilde{\sigma}_k^2)$: vector of k variances

$(\tilde{w}_1^2, \ldots, \tilde{w}_k^2)$: vector of k weights

Notation:

$\vec{S}$: vector containing the elements of set S in arbitrary order

$\vec{x} \circ \vec{y}$: concatenation of vectors, i.e. for $\vec{x} = (x_1, \ldots, x_n)$ and $\vec{y} = (y_1, \ldots, y_m)$, $\vec{x} \circ \vec{y} = (x_1, \ldots, x_n, y_1, \ldots, y_m)$

Output: $\theta = \{(w_i, \mu_i, \sigma_i)\}$ containing at most k tuples of mean and variance

if $l = 0$ **then**
  return $\vec{P}$;
else
  if $l \geq |R|$ **then**
    return $\theta = \{(\mu_i, \sigma_i)\}_i$ where $\vec{\mu} \circ \vec{R} = (\mu_i)_i$;
  else
    /* *sampling phase* */;
    sample a multiset S of size $\frac{1}{\alpha\varepsilon\delta}$ from R;
    $T \leftarrow \left\{\mu(S') | S' \subset S, |S'| = \frac{1}{\varepsilon\delta}\right\}$;
    $\mathcal{M}_k \leftarrow \emptyset$;
    for $t \in T$ **do**
      $\mathcal{M}_k \leftarrow \mathcal{M}_k \cup$ APPROX-MEANS$(R, l-1, \{\vec{\mu} \circ (t) | \vec{\mu} \in \mathcal{M}_{k-l}\}, \Sigma)$;
    end
    /* *pruning phase* */;
    $N \leftarrow$ set of $\frac{|R|}{2}$ points x from R with smallest minimum negative complete-data log-likelihood cost wrt. the weighted component given by $(\tilde{w}_i, \mu_i, \tilde{\sigma}_i^2)$ for $i \in [j]$, i.e.

$$\min_{i \in [j]} \left\{ \frac{d}{2} \ln(2\pi\tilde{\sigma}_i^2) + \frac{1}{2\tilde{\sigma}_i^2} \|x - \mu_i\|^2 - \ln(\tilde{w}_i) \right\}$$

    $\mathcal{M}_k \leftarrow \mathcal{M}_k \cup$ APPROX-MEANS$(R \setminus N, l, \mathcal{M}_{k-l}, \Sigma)$;
    return the candidate $\theta = \{(w_i, \mu_i, \sigma_i)\}_i$, $(\mu_i) \in \mathcal{M}_k$, which has minimal cost $\mathcal{L}_X(\theta)$;
  end
end

4.2 Applying the ABS Algorithm

In the following we analyze Algorithm 2. We show that the algorithm can be used to construct means such that, together with appropriate approximations of the weights and variances, we obtain a CMLE solution with costs close to the costs of the given CMLE solution.

Theorem 26. *Let $\tilde{\sigma}_i \in [\sigma_i^2, (\sigma_i^2)^{(1+\varepsilon)}]$ and $\tilde{w}_k \geq \frac{1}{(1+\varepsilon)} w_k$ for $i \in [k]$. Algorithm 2 started with $(X, k, \emptyset, (\tilde{\sigma}_1^2, \ldots, \tilde{\sigma}_k^2))$ computes a tuple $(\tilde{\mu}_1, \ldots, \tilde{\mu}_k)$ such that with probability at least $\left(\frac{1-\delta}{5}\right)^k$*

$$\mathcal{L}_X((\tilde{w}_i, \tilde{\mu}_i, \tilde{\sigma}_i^2)_{i\in[k]}) \leq (1+\varepsilon)\mathcal{L}(X) \ .$$

The running time of the algorithm is bounded by $|X|\, d\, 2^{\mathcal{O}(k/\varepsilon \cdot \log(k/\varepsilon^2))}$.

Let $\dot{\bigcup}_{i=1}^{k} C_i$ be a partition of X into optimal CMLE clusters. We introduce

$$C_{[i,j]} = \dot{\bigcup}_{t=i}^{j} C_t$$

as a short notation for the disjoint union of clusters i through j. We assume that the C_i are numbered by the order their approximate means $\tilde{\mu}_i$ are found by the superset-sampling technique.

Now, let $X = R_0 \supseteq R_i \supseteq \cdots \supseteq R_{k-1}$ be a sequence of input sets computed by the algorithm, such that

$$|C_i \cap R_{i-1}| \geq \alpha |R_{i-1}|.$$

Without loss of generality assume that each R_i is the largest of these sets with this property.

By using Lemma 17, we obtain the following Lemma.

Lemma 27 (By Superset-Sampling)**.** *With probability at least $((1-\delta)/5)^k$ we have*

$$\|\tilde{\mu}_i - \mu(C_i \cap R_{i-1})\|^2 \leq \frac{\varepsilon}{|C_i \cap R_{i-1}|} \sum_{x \in C_i \cap R_{i-1}} \|x - \mu(C_i \cap R_{i-1})\|^2$$

for all $i \in [K]$.

By $N_i \coloneqq R_{i-1} \setminus R_i$ we denote the set of points remove between two sampling phases. Using these definitions we can see that

$$\dot{\bigcup}_{i=1}^{k} (C_i \cap R_{i-1}) \ \dot{\cup} \dot{\bigcup}_{i=1}^{k} \left(C_{[i+1,k]} \cap N_i\right)$$

is a disjoint partition of X. Each set $C_i \cap R_{i-1}$ on the left side contains the points that the mean $\tilde{\mu}_i$ has been sampled from. The sets $C_{[i+1,k]} \cap N_i$ on the right side contain points incorrectly assigned to $\{\tilde{\mu}_1, \ldots, \tilde{\mu}_i\}$ during the pruning phases between the sampling of $\tilde{\mu}_i$ and $\tilde{\mu}_{i+1}$.

Denote by θ_i the parameters of the first i weighted Gaussians obtained by the algorithm, i.e.

$$\tilde{\theta}_i = ((\tilde{w}_1, \tilde{\mu}_1, \tilde{\sigma}_1), \ldots, (\tilde{w}_i, \tilde{\mu}_i, \tilde{\sigma}_i)) \ .$$

Lemma 28 (cf. Claim 4.8 in [Ack09])**.**

$$\mathcal{L}_{C_{[i+1,k]} \cap N_i}(\tilde{\theta}_i) \leq 8\alpha k \mathcal{L}_{C_{[1,i]} \cap R_{i-1}}(\tilde{\theta}_i)$$

Proof. As in [Ack09, p. 70ff], with "cost" replaced by "$\mathcal{L}$". □

Denote by $\mathrm{cost}(P, C)$ the k-means cost of a point set P wrt. a set of means C.

Lemma 29 (cf. Claim 4.9 in [Ack09])**.** *For every $i \in [k]$ we have*

$$\mathrm{cost}(C_i \cap R_{i-1}, \tilde{\mu}_i) \leq (1+\varepsilon)\,\mathrm{cost}(C_i, \mu_i) \ .$$

Proof. As in [Ack09, p. 70ff], using that optimal means in CMLE are means of the optimal CMLE clusters. □

Given appropriate approximate variances, we can conclude that a similar bound holds wrt. the complete-data log-likelihood.

Lemma 30. *Given $\tilde{\sigma}_i \in [\sigma_i^2, (\sigma_i^2)^{(1+\varepsilon)}]$ and $\tilde{w}_i = \frac{n_i}{|X|}$ with $n_i \in [|C_i|, (1+\varepsilon)|C_i|]$, we have*

$$\mathcal{L}_{C_i \cap R_{i-1}}(\tilde{w}_i, \tilde{\mu}_i, \tilde{\sigma}_i^2) \leq (1+\varepsilon)\mathcal{L}_{C_i}(w_i, \mu_i, \sigma_i^2) \ .$$

Proof.

$$\mathcal{L}_{C_i \cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) = \frac{|C_i \cap R_{i-1}|d}{2}\ln(2\pi\tilde{\sigma}_i^2) + \frac{1}{2\tilde{\sigma}_i^2}\,\mathrm{cost}(C_i \cap R_{i-1}, \tilde{\mu}_i) - |C_i \cap R_{i-1}|\ln(\tilde{w}_i) \ .$$

We have

$$\ln(2\pi\tilde{\sigma}_i^2) \leq \ln(2\pi(\sigma_i^2)^{(1+\varepsilon)}) = (1+\varepsilon)\ln(2\pi\sigma_i^2) \ .$$

Furthermore, Using that $|C_l| \leq n_l \leq (1+\varepsilon)|C_l|$ for all $l = 1, \ldots, K$, we obtain $\tilde{w}_k \geq \frac{|C_k|}{|X|}$. Hence,

$$\begin{aligned} -\ln(\tilde{w}_i) \cdot |C_i \cap R_{i-1}| &\leq -\ln(\tilde{w}_i) \cdot |C_i| \\ &\leq -\ln\left(\frac{|C_i|}{|X|}\right)|C_i| && \text{(by Equation (3))} \\ &= -\ln(w_i) \cdot |C_i| \end{aligned}$$

By Lemma 29 and $\tilde{\sigma}_i^2 \geq \sigma_i^2$,

$$\frac{1}{2\tilde{\sigma}_i^2}\,\mathrm{cost}(C_i \cap R_{i-1}, \tilde{\mu}_i) \leq (1+\varepsilon)\frac{1}{2\sigma_i^2}\,\mathrm{cost}(C_i, \mu_i) \ .$$

From this and by using that $\sigma_i^2 = \frac{1}{|C_i|d}\,\mathrm{cost}(C_i, \mu_i)$, we conclude

$$\begin{aligned} \mathcal{L}_{C_i \cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) &\leq (1+\varepsilon)\frac{|C_i|d}{2}\ln(2\pi\sigma_i^2) + (1+\varepsilon)\frac{1}{2\sigma_i^2}\,\mathrm{cost}(C_i, \mu_i) - \ln(w_i)|C_i| \\ &\leq (1+2\varepsilon)\mathcal{N}_1(C_i) - \ln(w_i)|C_i| \\ &\leq (1+2\varepsilon)\mathcal{L}_{C_i}(\mu_i, \sigma_i^2) \ . \end{aligned}$$

Running Algorithm 2 with $\varepsilon/3$ instead of ε yields the claim. □

Analogously to [Ack09], we can prove Theorem 26 as follows.

Proof of Theorem 26. Let $\tilde{\theta}_k = (\tilde{\mu}_i, \tilde{\sigma}_i^2)_{i\in[k]}$. Then,

$$\begin{aligned}
\mathcal{L}_X(\tilde{\theta}_k) &\le \sum_{i=1}^{k} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) + \sum_{i=1}^{k-1} \mathcal{L}_{C_{[i+1,k]}\cap N_i}(\tilde{\theta}_k) \\
&\le \sum_{i=1}^{k} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) + 8\alpha k \sum_{i=1}^{k-1} \mathcal{L}_{C_{[1,i]}\cap R_{i-1}}(\tilde{\theta}_k) && \text{(due to Lemma 28)} \\
&\le \sum_{i=1}^{k} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) + 8\alpha k \sum_{i=1}^{k-1}\sum_{t=1}^{i} \mathcal{L}_{C_t\cap R_{i-1}}(\tilde{\mu}_t, \tilde{\sigma}_t^2) \ .
\end{aligned}$$

Since $R_i \subseteq R_{i-1}$, we have $C_t \cap R_{i-1} \subseteq C_t \cap R_{t-1}$. Hence,

$$\begin{aligned}
\sum_{i=1}^{k-1}\sum_{t=1}^{i} \mathcal{L}_{C_t\cap R_{i-1}}(\tilde{\mu}_t, \tilde{\sigma}_t^2) &\le \sum_{i=1}^{k-1}\sum_{t=1}^{i} \mathcal{L}_{C_t\cap R_{t-1}}(\tilde{\mu}_t, \tilde{\sigma}_t^2) \\
&\le k \sum_{i=1}^{k-1} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) \ .
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathcal{L}_X(\tilde{\theta}_k) &\le \sum_{i=1}^{k} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) + 8\alpha k^2 \sum_{i=1}^{k-1} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) \\
&\le (1 + 8\alpha k^2) \sum_{i=1}^{k} \mathcal{L}_{C_i\cap R_{i-1}}(\tilde{\mu}_i, \tilde{\sigma}_i^2) \\
&\le (1 + 8\alpha k^2)(1+\varepsilon)\mathcal{L}(X) \ . && \text{(by Lemma 30)}
\end{aligned}$$

Finally, running the algorithm for $\varepsilon := \varepsilon/2$ and $\alpha = \theta(\varepsilon/k^2)$ yields the theorem.

□

5 Special Cases

5.1 Weighted K-Means (Identical Covariances)

In this section we consider a restricted version of the CMLE problem where we are only interested in Gaussian mixture models where all components share the same fixed spherical covariance matrix, i.e. parameters $\theta = \{(w_k, \mu_k, \Sigma_k)\}_{k\in[K]}$ where $\Sigma_k = \frac{1}{2\beta} I_d$ for all $k \in [K]$. We call this problem the *Weighted K-Means* (WKM) problem.

Problem 31 (WKM)**.** *Given a finite set $X \subset \mathbb{R}^d$ and an integer $K \in \mathbb{N}$, find a partition $\mathcal{C} = \{C_1, \ldots, C_K\}$ of X into K disjoint subsets and K weighted means $\theta = \{(w_k, \mu_k)\}_{k=1}^K$, where $\mu_k \in \mathbb{R}^D$, $w_k \in \mathbb{R}$, and $\sum_{k=1}^K w_k = 1$, minimizing*

$$\mathcal{L}_X^{wm}(\theta, \mathcal{C}) = \sum_{k=1}^{K} \beta \left(\sum_{x\in C_k} \|x - \mu_k\|^2 \right) - \ln(w_k) \cdot |C_k| \ .$$

We denote the minimal value by $OPT_{wm}(X, K)$.

Corollary 32. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$, and $\delta, \varepsilon > 0$. Let $X = \dot{\bigcup}_{k=1}^{K} C_k$ be a well-defined solution for the WKM problem. There is an algorithm that computes K weighted means $\theta = \{(\tilde{w}_k, \tilde{\mu}_k)\}_{k=1}^{K}$ such that with probability at least $1 - \delta$*

$$\mathcal{L}_X^{wm}((\tilde{w}_i, \tilde{\mu}_i)_{i\in[K]}) \leq (1+\varepsilon) OPT_{wm}(X, K) \ .$$

The running time of the algorithm is bounded by

$$|X|\, d\, 2^{\mathcal{O}(K/\varepsilon \cdot \log(K/\varepsilon^2))} \cdot (\log(f(K)))^K \ .$$

Proof. Use a grid search to obtain candidates for the weights, then apply the ABS algorithm. □

5.2 Uniform Weights

In this section we consider a restricted version of the CMLE problem where we are only interested in Gaussian mixture models with fixed uniform weights, i.e. parameters $\theta = \{(w_k, \mu_k, \Sigma_k)\}_{k\in[K]}$ where $w_k = 1/K$ for all $k \in [K]$. We denote this problem by *Uniform Complete-Data Maximum Likelihood Estimation* (UCMLE).

Problem 33 (UCMLE)**.** *Given a finite set $X \subset \mathbb{R}^d$ and an integer $K \in \mathbb{N}$, find a partition $\mathcal{C} = \{C_1, \ldots, C_K\}$ of X into K disjoint subsets and K spherical Gaussians with parameters $\theta = \{(\mu_k, \sigma_k^2)\}_{k=1}^{K}$ minimizing*

$$\begin{aligned}\mathcal{L}_X^{unif}(\theta, \mathcal{C}) &= \sum_{k=1}^{K} \mathcal{L}_{C_k}(\mu_k, \sigma_k^2) \\ &= \sum_{k=1}^{K} \frac{|C_k| d}{2} \ln(2\pi\sigma_k^2) + \frac{1}{2\sigma_k^2} \left(\sum_{x \in C_k} \|x - \mu_k\|^2 \right) \ .\end{aligned}$$

We denote the minimal value by $OPT_{unif}(X, K)$.

Corollary 34. *Let $X \subset \mathbb{R}^d$, $K \in \mathbb{N}$, and $\delta, \varepsilon > 0$. Let $X = \dot{\bigcup}_{k=1}^{K} C_k$ be a well-defined solution for the UCMLE problem. There is an algorithm that computes K spherical Gaussians $\theta = \{(\tilde{\mu}_k, \tilde{\sigma}_k^2)\}_{k=1}^{K}$ such that with probability at least $1 - \delta$*

$$\mathcal{L}_X^{unif}((\tilde{\mu}_i, \tilde{\sigma}_i^2)_{i\in[K]}) \leq (1+\varepsilon) OPT_{unif}(X, K) \ .$$

The running time of the algorithm is bounded by

$$|X|\, d\, \log(1/\delta)\, 2^{\mathcal{O}(K/\varepsilon \cdot \log(K/\varepsilon^2))} \left(\log(\log(\Delta^2)) + 1\right)^K \ ,$$

where $\Delta^2 = \max_{x,y\in X}\{\|x - y\|^2\}$.

Proof. Use a grid search to obtain candidates for the variances, then apply the ABS algorithm. □

References

[ABS10] Marcel R. Ackermann, Johannes Blömer, and Christian Sohler. Clustering for metric and nonmetric distance measures. *ACM Trans. Algorithms*, 6(4):59:1–59:26, September 2010.

[Ack09] Marcel R. Ackermann. *Algorithms for the Bregman k-Median Problem*. PhD thesis, University of Paderborn, 2009.

[CG92] Celeux and Govaert. A Classification EM Algorithm for Clustering and Two Stochastic Versions. *Comput. Stat. Data Anal.*, 14(3), 1992.

[IKI94] M. Inaba, N. Katoh, and H. Imai. Applications of Weighted Voronoi Diagrams and Randomization to Variance-based K-clustering. In *Proceedings of the Tenth Annual Symposium on Computational Geometry*, SoCG '94, pages 332–339, New York, NY, USA, 1994. ACM.